Journal of Computer Technology and Software

ISSN:2998-2383

Vol. 3, No. 1, 2024

Corporate Cash Flow Forecasting Based on Counterfactual Time Series and Strategy Simulation

Yunfei Nie
Brandeis University, Waltham, USA
samnie519@gmail.com

Abstract: This study addresses the instability and limited decision usability of corporate cash flow forecasting under conditions of multi-source heterogeneous data coupling, pronounced inter-period lag transmission, and frequent strategic interventions. A unified modeling approach based on counterfactual time series is proposed. The method organizes corporate operating state features, controllable action variables, and exogenous environmental factors into structured historical sequences. An encoder is used to learn low-dimensional latent state representations. An explicit state transition mechanism then models the dynamic evolution of latent states driven by actions and environments. A decoder generates multi-step cash flow forecasts. To support counterfactual analysis, alternative action sequences are substituted under the same historical starting point to construct counterfactual paths. Comparable cash flow trajectories are produced to characterize fund response differences under different strategies. In addition, inflow and outflow components are modeled separately to ensure consistency with financial definitions. Representation alignment and counterfactual consistency constraints are introduced to suppress the influence of selection bias and noise-driven correlations on inference reliability. Sensitivity analyses are further conducted with respect to the number of attention heads, input noise injection intensity, proportion of extreme fluctuation samples, and training data ratio. These analyses systematically characterize the impact boundaries of key configurations and data conditions on error metrics. The framework, therefore, provides an interpretable and scenario-driven modeling and evaluation basis for corporate cash flow forecasting and cash management.

Keywords: Counterfactual analysis; cash flow time-series modeling; robustness assessment; contextualized cash management

1. Introduction

In the context of intensified digital operations and parallel reinforcement of risk constraints, corporate cash flow forecasting is evolving from a traditional financial management tool into a core capability that permeates operational decision making, resource allocation, and risk early warning. Cash flow directly reflects actual fund inflows and outflows within a given period [1]. It captures a firm's payment capacity and liquidity conditions. It affects the stability of daily operations and shapes financing arrangements, investment pacing, supply chain coordination, and budget execution. Compared with profit-based accounting indicators, cash flow is closer to real payment and collection processes. It is more sensitive to external shocks, internal management efficiency, and changes in business structure. Under environments with intensified order volatility, shortened credit cycles, and rapid transmission of interest rates and costs, the accuracy and interpretability of cash flow forecasting have become critical prerequisites for enhancing corporate resilience, reducing liquidity disruption risk, and achieving stable growth.

However, corporate cash flow forecasting has long faced compound challenges characterized by strong temporal dependence, strong heterogeneity, and strong causal interference. Cash flow is jointly driven by multiple sources, including sales collections, procurement payments, tax obligations, payroll expenses, and capital expenditures. This leads to lag effects, rhythmic fluctuations, and structural breaks. At the same time, firms differ substantially in industry attributes, customer composition, payment terms, financing channels, and governance mechanisms. As a result, the same indicator may carry different economic meanings and influence paths across firms. In real business processes, there are also many strategic behaviors and managerial interventions [2]. Examples include adjustments to credit policies, renegotiation of payment terms, price promotions, inventory replenishment, receivables collection efforts, and the advancement or postponement of payments. These interventions introduce non-stationarity into the data. Models that rely only on correlation learning may mistake policy-driven patterns for stable regularities. They can therefore fail when the environment shifts or strategies change.

From a managerial decision perspective, firms are not only concerned with how much future cash flow will be generated. They also care about why it will take that form and how it would change under alternative actions. This gives cash flow forecasting an inherent decision support role. Forecast outputs are expected to inform fund planning, financing limits, working capital management, and risk mitigation strategies. Point forecasts alone are often insufficient for high-quality decisions [3]. Cash flow management involves an interplay between controllable and uncontrollable factors. Managers need to evaluate the impact of boundaries and the potential costs of alternative strategies. For instance, rising collection pressure may prompt adjustments to credit terms, stronger collection efforts, or changes in promotional intensity. Cost increases may lead to early price locking or delayed procurement. Financing constraints may require reduced capital expenditure or optimized inventory turnover. If a forecasting model cannot distinguish managerial interventions from external demand changes, its outputs are difficult to translate into actionable strategies. Misattributed causal effects may even induce counterproductive decisions [4].

Against this background, introducing a counterfactual perspective to reformulate the cash flow forecasting problem is of substantial importance. Counterfactual time series emphasize the construction of comparative paths under different values of controllable decision variables, given historical information and business constraints. They aim to answer how future cash flows would change if alternative actions were taken. Unlike traditional approaches that focus only on mapping the past to the future, counterfactual modeling highlights causal relationships between interventions and outcomes. It provides more robust forecasting foundations in environments with multiple concurrent strategies and frequent rule adjustments. More importantly, counterfactual sequences enable scenario analysis for managers. Under the same economic conditions and business states, different strategy combinations can be compared through their projected cash flow trajectories. This shifts cash management from experience-driven judgment to evidence-driven reasoning. It improves the forward-looking consistency of budgeting, fund allocation, and risk control [5].

As corporate governance and external regulation increasingly emphasize transparency, interpretability, and accountability for risk, counterfactual time series-based cash flow forecasting offers broader practical value. First, it supports the construction of early warning and stress testing systems centered on cash flow. It helps assess funding gaps and response strategies under extreme scenarios. Second, it facilitates refined working capital management. Through counterfactual evaluation of collection and payment strategies, it enhances capital turnover efficiency and supply chain coordination. Third, it contributes to financing communication and credit assessment by improving the credibility and external acceptability of forecasts through clearer mechanism explanations. Overall, research on counterfactual time series-based corporate cash flow forecasting not only improves predictive accuracy and cross-environment robustness. More importantly, it connects forecasts with controllable decision variables. It forms an interpretable, scenario-based, and actionable cash management support framework. This provides methodological support for firms to achieve stable operations and controllable risk under uncertainty.

2. Related Work

Existing studies on corporate cash flow forecasting can be broadly categorized into two main streams, namely statistical modeling and machine learning based approaches. Statistical methods usually start from time series decomposition and linear or nonlinear regression frameworks. They describe cash flow dynamics through trend components, seasonal components, lag terms, and error structures. Exogenous variables are often incorporated to enhance sensitivity to macroeconomic and operational environment changes. These methods feature clear structures and relatively strong interpretability. They are also more suitable for small sample settings and can be easily aligned with budgeting systems and accounting standards. However, their ability to capture non-stationary shocks, structural breaks, long-range dependencies, and multi-source interactions is limited. Strong prior assumptions or frequent manual parameter tuning are often required. When strategic interventions, payment term adjustments, or collection structure changes occur, models relying on stationarity or fixed parameters tend to suffer from extrapolation bias. They also struggle to distinguish between differences driven by strategic changes and those caused by environmental shifts [6].

Machine learning based research has progressively strengthened the capacity to learn complex nonlinear relationships and high-dimensional feature interactions. Common practices include feature-engineered regression and ensemble learning methods, as well as deep temporal models that emphasize end-to-end representation learning. The former typically encode financial statement indicators, transaction records, receivables and payables structures, and operational metrics into feature vectors, and then learn cash flow mappings through nonlinear models. The latter model sequences directly. They rely on recurrent structures, convolutional structures, or attention mechanisms to capture short-term fluctuations and long-term dependencies. They can also learn dynamic coupling across multiple variables. These approaches often show greater potential in predictive accuracy, especially for large-scale data and complex business settings. At the same time, they face two major challenges. One is limited interpretability, which makes it difficult to align predictions with financial logic and business actions. The other is sensitivity to distribution shifts and strategy changes. Models may learn patterns that are correlated but not causal. This can lead to degraded stability under policy switches, payment term adjustments, or business restructuring.

To address demands for interpretability and transferability in cash flow forecasting, some studies have introduced structured priors and multi-source relational modeling. One line of work decomposes cash flow into sub-processes such as collections, payments, taxes, financing, and investment. These components are forecast separately and then aggregated. This enhances economic meaning and traceability. Another line of work uses graph structures or hierarchical representations to model relationships among firms, customers, suppliers, contracts, and projects. Transaction networks and entity relations are incorporated as complementary information. This helps mitigate the limitations of pure sequence modeling in capturing implicit dependencies. In addition, some studies focus on uncertainty estimation and risk-sensitive forecasting. Interval or quantile predictions are provided to support cash buffer management and better reflect real decision scenarios. Although these directions improve interpretability and robustness, most methods still focus on making forecasts more reliable. They offer limited answers to how cash flow responds when controllable strategy variables change. Systematic characterization of the mechanisms linking interventions and outcomes remains insufficient [7].

In recent years, causal inference and counterfactual learning have provided new theoretical foundations and methodological tools for time series forecasting. Their core objective is to distinguish observed correlations from intervention effects and to enable scenario evaluation through counterfactual sequence construction. In corporate cash flow contexts, credit policies, payment schedules, inventory replenishment, promotion intensity, and collection efforts can all be viewed as controllable interventions. If a model can identify the marginal effects of these interventions under temporal dependence and external environmental influences, forecasting can be elevated to decision support. However, applying counterfactual ideas to corporate cash flow remains challenging. Interventions are often high-dimensional and time varying. They involve delayed effects and feedback loops. Observational data suffer from selection bias and confounding, with interventions interacting with operational states. Strong heterogeneity across firms further leads to inconsistent effects of the same intervention. As a result, recent research explores combining representation learning with causal constraints. This includes learning transferable latent states, explicitly modeling confounding and time-varying effects, and embedding business constraints into sequence generation. These efforts aim to construct counterfactual time series that better reflect operational logic. This direction lays an important foundation for shifting corporate cash flow forecasting from numerical fitting toward mechanism-based and scenario-driven prediction.

3. Proposed Framework

In corporate cash flow forecasting, observed operating and financial information is represented as a multivariate time series, and a controllable decision sequence is explicitly introduced to support counterfactual inference. Its overall model architecture is shown in Figure 1.

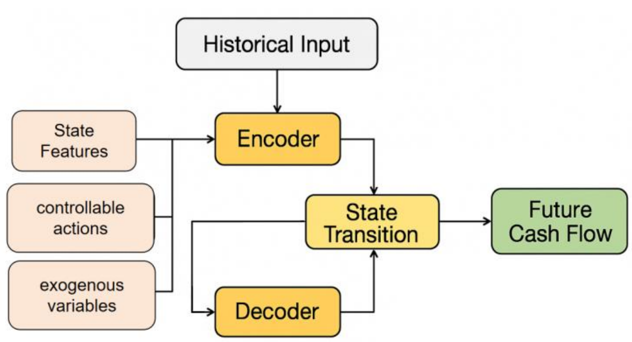

Figure 1. Overall model architecture

Given a historical window of length T, let the state feature be x_t, the controllable decision be a_t, the exogenous environment be c_t, and the cash flow be y_t. The historical input can be written as:

$$H_T = \{(x_t, a_t, c_t, y_t)\}_{t=1}^{T}$$

The goal is to predict future cash flows in step H, given the condition H_T, while retaining the ability to extrapolate scenarios for different decision paths.

$$y_{T+1:T+H} = F_\theta(H_T),\ \ y_{T+1:T+H} \in R^H)$$

To characterize the temporal dependence and multi-source coupling of cash flows, the method uses latent state representation to compress historical information and drives the evolution of latent states with controllable decision-making. First, observations from the last L steps are concatenated into the input segment s_t, and the latent state is obtained through an encoder:

$$z_t = f_\theta(s_t)$$

Then, the decision term and exogenous term are explicitly introduced using a state transition function to obtain the next potential state:

$$z_{t+1} = g_\theta(z_t, a_t, c_t)$$

Finally, the decoder outputs the predicted cash flow value:

$$\hat{y}_{t+1} = h_\theta(z_{t+1})$$

This structure separates "representation learning of historical information" from "decision-driven dynamic evolution," enabling the model to capture both long-term trends and short-term fluctuations, while also providing an interface for counterfactual interventions in form.

In counterfactual deduction, the decision sequence is treated as an intrusive variable. A counterfactual path is constructed for an alternative strategy under the same historical state, and the corresponding future cash flow trajectory is generated. Let represent external intervention in the decision-making process. Then, under the same historical starting point z_t^{cf}, the counterfactual potential state is recursively derived as follows:

$$z_{t+1}^{cf} = g_\theta(z_t^{cf}, a_t, c_t)$$

And based on this, a counterfactual cash flow forecast is obtained:

$$y_{t+1}^{cf} = h_\theta(z_{t+1}^{cf})$$

This training method enables the model to learn the time structure of cash flow and output comparable counterfactual time series in the form of "different decision paths under the same historical state", thereby providing a predictable quantitative basis for corporate funding plans, credit term adjustments, payment rhythm optimization, and other strategies.

4. Experimental Analysis

4.1 Dataset

This study adopts the Financial Statement Data Sets publicly released by the United States Securities regulator as the open source data foundation for corporate cash flow forecasting. The dataset extracts numerical values of financials from XBRL filings submitted by listed firms. It covers core statements, including the balance sheet, income statement, and cash flow statement. It allows direct access to cash flow items related to operating, investing, and financing activities. The data are naturally organized by firm and reporting period. This structure is suitable for firm-level cash flow forecasting and counterfactual scenario construction. The dataset is updated on a quarterly basis. Its time span extends from 2009 to the recent quarters. This supports long sequence modeling and robust analysis across economic cycles.

In terms of data organization, the Financial Statement Data Sets are provided as quarterly compressed packages. They mainly consist of four types of tables, namely SUB, NUM, TAG, and PRE. The SUB table identifies each submission and reporting entity, such as accession identifiers and firm identifiers. The NUM table records numerical facts for each reporting item. It includes the associated tag and the report period end date. The TAG table provides textual definitions and attributes of tags. The PRE table describes the presentation positions and statement affiliations of tags across different reports. This structure enables stable identification of cash flow statement items and consistent alignment across firms. Based on key relationships among these tables, the original item-level and statement-level facts can be reconstructed into structured firm-quarter time series data. This provides a reliable data interface for subsequent state encoding, decision intervention modeling, and counterfactual path generation.

Consistent with the theme of counterfactual time series-based corporate cash flow forecasting, the dataset can be directly used to construct prediction targets and multiple source inputs. Core cash flow items from the cash flow statement serve as target sequences, such as net cash flows from operating activities and their related components. Numerical items related to working capital, capital expenditures, and financing activities are extracted from the cash flow statement, balance sheet, and income statement as state features. Metadata in the SUB table, including industry classification, is used to form exogenous environment variables. This supports aligned modeling across firms and operating contexts. On this basis, financial items related to credit policies, collection and payment timing, and investment and financing arrangements can be further mapped to proxy variables for controllable decisions or actions. This provides data support for constructing counterfactual strategy sequences under the same historical conditions. It allows forecasting to focus not only on future values but also on scenario-based and actionable strategic implications.

4.2 Experimental Results

This article first presents the results of the comparative experiments, as shown in Table 1.

Table 1. Comparative experimental results

Method	MSE	MAE	MAPE	RMSE
LSTM [8]	0.0290	0.1270	0.098	0.1703
BI-LSTM [9]	0.0256	0.1180	0.091	0.1600
Transformer [10]	0.0217	0.1090	0.085	0.1473
Informer [11]	0.0199	0.1040	0.080	0.1411
ITransformer [12]	0.0178	0.0980	0.076	0.1334
TimeMixer [13]	0.0163	0.0940	0.073	0.1277
Ours	0.0134	0.0860	0.066	0.1158

As shown in Table 1, the models exhibit a clear performance hierarchy on the corporate cash flow forecasting task, which is characterized by strong nonstationarity and strong coupling. Traditional LSTM and BI LSTM models provide basic sequence memory. However, their errors remain relatively large under scenarios with multiple source drivers and lag effects. This indicates that relying only on gated recurrent structures is often insufficient to capture the complex dynamics formed by the joint effects of operating collections, payment schedules, and financing arrangements. When the modeling paradigm shifts from recurrent structures to attention-dominated architectures, overall errors decrease markedly. This suggests that global dependency modeling is more effective in handling inter-period transmission and long-cycle capital turnover structures.

Within the attention-based model family, the development paths of Transformer, Informer, iTransformer, and TimeMixer reflect two main trends. The first is a more efficient aggregation of long sequence information. The second is representation schemes that are more aligned with time series structures. Compared with the basic Transformer, Informer achieves further reductions in MSE and MAE. This indicates more effective utilization of long-range historical information. It is better suited to common cash flow patterns such as delayed collections and cross-period payment arrangements. The performance of iTransformer and TimeMixer continues to improve. In particular, the gradual decrease in MAPE shows that the models reduce not only average errors but also relative errors. This stability is important for cash flow variables whose scale varies significantly with firm size and operating cycles.

The proposed model achieves the best performance across all four metrics. MSE decreases from 0.0163 for TimeMixer to 0.0134. MAE decreases from 0.0940 to 0.0860. MAPE decreases from 0.073 to 0.066. RMSE decreases from 0.1277 to 0.1158. These improvements indicate stronger suppression of extreme fluctuations and tail risks. In the context of counterfactual time series modeling, such consistent gains suggest a clearer characterization of the structural relationships between controllable decision variables and cash flow responses. Under the same historical states, scenario projections along different strategy paths are therefore less affected by spurious correlations. Lower RMSE and MSE imply not only more accurate point forecasts but also fewer large deviation events. This has direct practical value for liquidity safety margins and cash flow risk management.

From an application perspective, the reduction in MAPE is particularly important for corporate cash flow management. Firms are often more sensitive to relative deviations when planning funds and arranging credit lines. This is especially true during periods of cash flow contraction or expansion. Relative errors can directly affect budget robustness and contingency funding preparation. The proposed model improves both relative and absolute errors at the same time. This indicates stronger adaptability across firms of different sizes and across different time periods. It supports transferable forecasting capability in environments with frequent strategy adjustments. For counterfactual modeling, this stability also implies that predicted trajectories under hypothetical action sequences are more likely to preserve reasonable monotonicity and economic consistency. This strengthens support for scenario analysis involving credit term adjustments, payment schedule optimization, and financing decisions.

This paper also presents the effect of the number of attention heads on the experimental results, and the experimental results are shown in Figure 2.

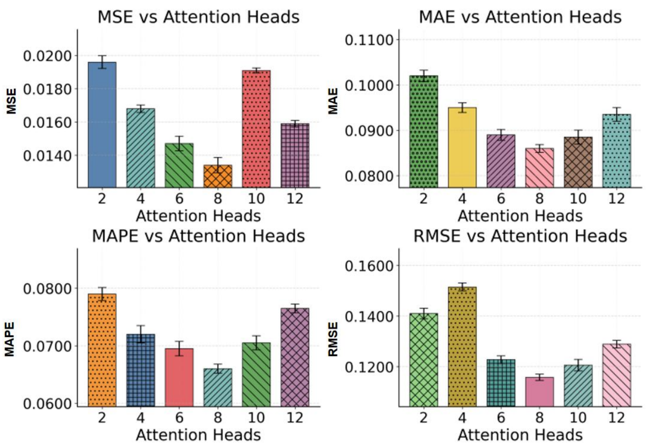

Figure 2. The effect of the number of attention heads on the experimental results

The four subplots in Figure 2 show that the number of attention heads has a significant impact on the cash flow forecasting performance of the proposed model. This impact follows a consistent structural pattern across all four error metrics. In general, as the number of heads increases from a small value to a moderate level, the model becomes better at capturing inter-period dependencies and multivariate coupling. Error metrics decrease steadily. This indicates that finer-grained attention decomposition helps learn dynamic transmission paths in cash flow that are jointly driven by delayed collections, payment schedules, financing activities, and operations. This behavior matches the strong nonlinearity and long lag characteristics of corporate cash flow series. It suggests that the number of heads is a key hyperparameter governing representation capacity and sequence modeling effectiveness.

For MSE and RMSE, errors decrease markedly as the number of heads increases from 4 to 8, with the lowest values observed at 8 heads. This indicates a better balance between overall fit and large deviation control at this setting. In cash flow forecasting, lower RMSE implies stronger suppression of extreme fluctuations and tail shocks. This is beneficial for liquidity safety margin management and early risk warning. When the number of heads further increases to 10 or 12, errors rise again. This suggests that an excessive number of heads may introduce redundant attention branches. Overly fine representation decomposition can make the model more prone to learning unstable correlation patterns under limited samples and noise. Generalization performance, therefore, deteriorates.

Changes in MAE and MAPE further reveal sensitivity at the relative error level. MAE reaches its minimum around 8 heads. MAPE shows a similar decreasing then increasing trend. This indicates that a moderate number of heads reduces absolute deviations and improves relative stability across firms of different sizes and across periods with varying cash flow magnitudes. In corporate cash flow applications, improvements in MAPE often imply more reliable budgeting, fund allocation, and credit arrangements in proportional terms. This is especially important during phases of cash flow contraction, where lower relative errors can reduce the risk of funding gaps caused by misjudged strategies. With too few heads, attention capacity is insufficient to align multiple source drivers and inter-period dependencies. With too many heads, overfitting and noise amplification can occur, leading to degraded relative errors.

When considered in light of the counterfactual time series modeling objective of this study, these results also imply a more important mechanism-level insight. Counterfactual scenario analysis requires the model to produce stable and interpretable response curves to different decision paths under the same historical state. Too few attention heads result in overly coarse representations of the relationship between decision variables and cash flow responses. Reliable scenario comparison is then difficult. Too many heads can make the model overly sensitive to fragmented patterns. Counterfactual trajectories may exhibit unnecessary fluctuations under small strategy adjustments. This reduces credibility and operational value. A moderate number of heads, therefore, provides a better balance between expressive power and robustness. It improves forecasting accuracy and supports the construction of stable counterfactual cash flow paths for strategy evaluation and decision support.

Input noise injection is a common method for verifying the robustness of cash flow forecasting models, used to simulate observation uncertainties caused by acquisition errors, caliber biases, and abnormal disturbances in real-world business operations. As noise intensity changes, the model's extraction of key time-series dependencies and multi-source coupled signals is affected to varying degrees, thus impacting the stability of error metrics. This experiment aims to characterize the performance boundaries of the proposed algorithm under different noise environments, providing a basis for robust configuration in subsequent deployments. The experimental results are shown in Figure 3.

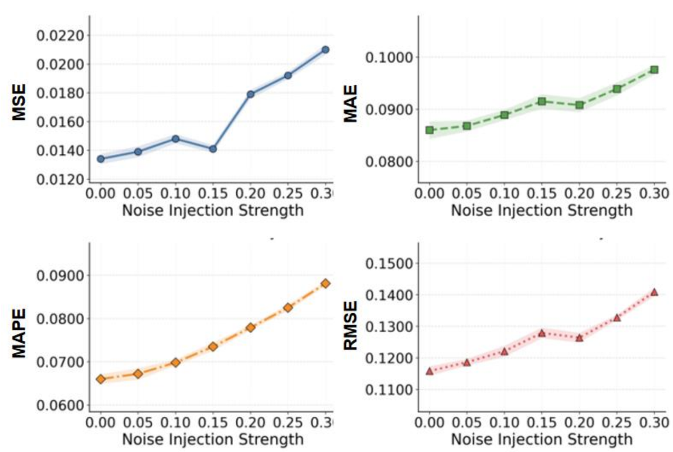

Figure 3. The effect of input noise injection intensity on experimental results

The overall trends in Figure 3 indicate that as the intensity of input noise injection increases, error levels across all four metricsoverall increasing trend with minor fluctuations. This shows that noise continuously reduces the effective signal density in sequence features. Cash flow forecasting becomes more dependent on corrupted historical collection and payment patterns. Cash flow series are typically driven by multiple sources and exhibit clear lag effects and inter-period transmission chains. Noise disrupts this cross-period consistency. It weakens the model's ability to align long-term dependencies. Errors, therefore, expand as perturbations increase. This behavior is consistent with real deployment environments, where data definition inconsistencies, collection delays, and abnormal entries undermine the stability of forecasting systems.

For MSE and RMSE, which place greater emphasis on large deviations, the curves display stronger amplification effects. The increase is more pronounced in high noise regimes. This reflects the fact that noise has a stronger impact on extreme errors and tail risks. In corporate cash flow applications, such large deviations often correspond to distorted estimates of funding gaps or underestimation of sudden liquidity pressure. The associated risk cost is typically higher than that of average errors. Changes in MSE and RMSE, therefore, indicate not only reduced predictive accuracy. They also signal weakened suppression of abnormal fluctuations and structural shocks under high disturbance conditions. Liquidity safety margins become more vulnerable. In practice, this highlights the importance of upstream data quality control and early anomaly detection.

The evolution of MAE and MAPE shows a relatively smoother increase. However, it still reflects the persistent impact of noise on stability. Rising MAE indicates larger absolute deviations. Rising MAPE indicates increasing relative errors. This means that even across different cash flow scales and firm sizes, proportional reliability declines as noise intensifies. In cash flow management contexts, changes in relative error directly affect the conservativeness of budgeting, credit limit setting, and short-term fund allocation. Higher noise levels require more robust buffering strategies or stricter threshold mechanisms. These measures help prevent systematic distortions in funding plans caused by proportional errors.

When interpreted together with the counterfactual time series objective of this study, these results reveal an important conclusion. The credibility of counterfactual scenario analysis depends strongly on the identifiability of structure in the input sequences. As noise increases, it becomes more difficult for the model to accurately capture the true effect pathways between controllable decision variables and cash flow responses. Differences across counterfactual sequences under strategy changes may be obscured by noise or driven by non-causal correlations. This weakens interpretability and operational value. Therefore, in decision support-oriented applications, noise sensitivity analysis should be a mandatory step before deployment. It helps determine acceptable ranges of data perturbation and robust configuration intervals. At the system level, it should be combined with input denoising, definition alignment, and anomaly correction mechanisms. These measures ensure stable outputs of counterfactual sequence generation and cash flow forecasting in complex operational environments.

Extreme cash flow fluctuations typically correspond to abnormal operational shocks or short-term cash flow pressures, and are key data factors for evaluating the risk robustness and tail error control capabilities of cash flow forecasting models. By adjusting the proportion of extreme fluctuation samples in the training data, the error sensitivity boundary of the algorithm in this paper can be systematically characterized under different tail strength conditions. The experimental results are shown in Figure 4.

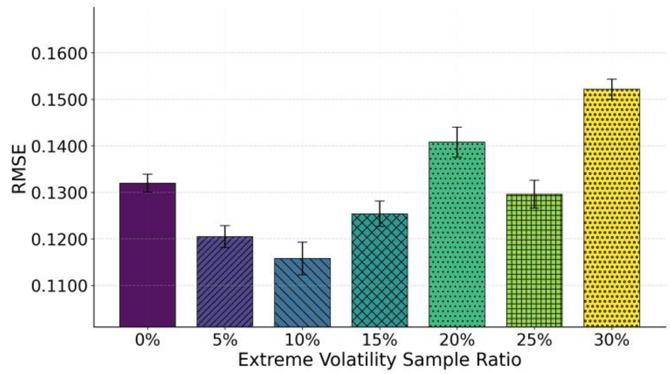

Figure 4. Data sensitivity analysis experiment on the impact of extreme cash flow volatility, sample proportion on cash flow forecast RMSE

Figure 4 shows that changes in the proportion of extreme cash flow fluctuation samples have a significant impact on RMSE and exhibit a non-monotonic sensitivity pattern. As the proportion increases from 0 percent to around 10 percent, RMSE decreases markedly. This indicates that introducing a moderate amount of tail fluctuation samples into the training set helps the model better learn cash flow dynamics under shock conditions. It improves adaptation to abnormal fluctuations. For corporate cash flow forecasting, this means that the model no longer fits only stable collection and payment patterns. It begins to capture sudden disturbances and inter-period transmission that commonly occur in real operations. Predictive reliability under risk scenarios is therefore enhanced.

When the proportion of extreme samples further increases to 15 percent and 20 percent, RMSE rises again. This suggests that an excessive number of tail samples alters the training distribution. The model becomes biased toward high volatility regions during representation learning. The fit to normal cash flow trajectories is then degraded. Cash flow sequences exhibit clear structural segmentation. The mechanisms governing fund turnover in normal regimes differ from those driving payment and collection behavior during shocks. When the tail proportion is too high, the model tends to overgeneralize shock-related features into normal state prediction. Overall root mean square error is amplified. This phenomenon reflects an inherent data-level trade-off. Tail samples are necessary for improving risk robustness. They can also become a source of distribution shift and overfitting.

At 25 percent, RMSE shows a partial decline. It then increases again at 30 percent. This further confirms that the sensitivity is not a simple linear relationship. It is more consistent with performance boundary shifts driven by mixture proportions of different regimes. For the counterfactual time series modeling objective of this study, such non-monotonicity has important implications. Counterfactual scenario analysis requires stable and mechanism-consistent responses across different state regimes. When the proportion of extreme samples is too low, counterfactual paths may underestimate cash flow risk under shock scenarios. When the proportion is too high, counterfactual responses may become overly sensitive to small strategy adjustments. This can introduce unnecessary volatility or amplify tail effects. Interpretability and operational value of scenario comparison are then reduced.

Therefore, this experiment points to a practical deployment insight. During data construction, a reasonable proportion range should be set for extreme fluctuation samples. This allows the model to learn key patterns of tail shocks without being dominated by tail distributions at the cost of normal regime accuracy. For corporate fund management, this corresponds to balancing risk coverage and baseline precision. The model can then support daily fund planning and cash flow scheduling while also providing credible counterfactual analysis under stress scenarios. By further combining strategies such as stratified sampling, tail weighting, or state partition-based modeling, this sensitivity can be mitigated without changing the model architecture. Robustness of counterfactual cash flow forecasting across multiple scenarios can thus be improved.

Finally, this paper presents the impact of the proportion of training data on the experimental results, which are shown in Table 2.

Table 2. The impact of the proportion of training data on experimental results

Training data percentage	MSE	MAE	MAPE	RMSE
50%	0.0219	0.1118	0.089	0.1480
55%	0.0197	0.1062	0.083	0.1403
60%	0.0178	0.1010	0.079	0.1334
65%	0.0172	0.0995	0.078	0.1311
70%	0.0158	0.0948	0.072	0.1257
75%	0.0146	0.0906	0.069	0.1208
80%	0.0134	0.0860	0.066	0.1158

Table 2 shows that increasing the training proportion from 50% to 80% leads to a steady reduction in all four error metrics, indicating that the proposed model benefits noticeably from larger training sets and is sensitive to data coverage. Corporate cash flow is jointly influenced by operating receipts, payment schedules, financing activities, and external conditions, which creates heterogeneous patterns with inter-period lags. With limited training data, the model may learn overly local regularities and become less robust around structural changes. As more training samples are provided, the model is exposed to a broader range of operating states and shock patterns, resulting in more consistent error reduction.

Across the training-ratio sweep, the absolute drops in MSE and RMSE are relatively large (MSE: 0.0219→0.0134; RMSE: 0.1480→0.1158). This suggests that expanding the data scale is particularly helpful for reducing the impact of rare, high-deviation cases. In practice, large forecast deviations often occur around collection delays, concentrated payments, or shifts in financing rhythm, which can induce tail fluctuations. If such events are under-represented during training, prediction errors can be amplified when similar states appear in evaluation. With higher training proportions, the model captures these patterns more adequately and constrains large-error events, supporting more reliable estimation of extreme funding gaps for liquidity buffers and early warning.

The concurrent decline in MAE and MAPE indicates improvements in both absolute and relative accuracy. Lower MAE reduces average deviation, which is beneficial for routine fund planning and rolling budgets. The drop in MAPE from 0.089 to 0.066 improves comparability across firms of different scales and across periods with different cash flow magnitudes. For the counterfactual time-series setting in this study, improved relative accuracy also reduces the risk that scenario trajectories under different strategy paths are distorted by scale-related error bias, thereby improving the interpretability and decision usefulness of counterfactual comparisons.

Finally, the values in Table 2 are obtained under a specific split; repeating the split can yield small variations due to randomness in sample composition, which is consistent with the heterogeneity and long-tail nature of cash flow data. For experimental reporting and deployment validation, fixed splits or repeated resampling should be used to quantify stability. Targeted inclusion of high-volatility periods can further improve coverage of rare but high-impact states, strengthening the robustness of subsequent counterfactual trajectory comparisons.

5. Conclusion

This study addresses the challenges of strong temporal dependence, multi-source heterogeneous coupling, and frequent strategic interventions in corporate cash flow forecasting under real operating conditions. It proposes a unified modeling framework oriented toward counterfactual time series. Cash flow forecasting is elevated from pure numerical extrapolation to a decision-oriented inference task. By explicitly modeling dynamic relationships among historical states, controllable actions, and exogenous environments, the approach achieves finer finer-grained representation of inter-period transmission and long lag effects at the prediction level. At the mechanism level, it provides a clear interface for generating cash flow trajectories under alternative strategy paths. As a result, forecasts are no longer limited to single-point values. They present scenario-based responses of inflows and outflows to strategy changes. This aligns more closely with practical needs in fund planning, working capital management, and risk mitigation.

From an application perspective, the primary contribution lies in offering more robust, interpretable, and actionable quantitative support for corporate cash flow management. Forecasting is commonly used across budgeting, fund allocation, credit term setting, payment schedule optimization, financing arrangement, and stress testing. Any prediction error can be amplified into funding gaps, mismatches, or risk control misjudgments. The use of counterfactual time series allows firms to evaluate potential consequences of different actions under the same historical and environmental constraints. It enables more comparative strategy assessment and reduces uncertainty associated with experience-driven decisions. In related domains such as auditing, financial shared services, and supply chain finance, the framework can serve as a general analytical layer. It connects business actions with financial outcomes. It supports the identification of high-risk fund paths and causes of abnormal fluctuations. Timeliness and targeting of risk warning and intervention can be improved.

The study also provides several extensible directions for future research and industrial deployment. Counterfactual modeling can be extended from single action sequences to multi-agent interaction settings. Examples include coordinated optimization of customer collection strategies and supplier payment strategies. More refined constraints can be introduced to reflect realistic boundaries such as payment terms, contract clauses, credit limits, and regulatory requirements. This would make generated scenarios closer to executable strategies. At the model level, adaptability to distribution shifts and sudden events can be enhanced. Transferable representations can be developed for macro shocks, industry cycle transitions, and policy changes. Interval-based or risk-stratified outputs can be further provided. These would offer quantifiable safety margins and resource buffer suggestions. With online updating and continual learning mechanisms, rolling forecasting and timely calibration can be supported. System stability can be maintained after data definition changes or strategy adjustments.

Overall, the proposed counterfactual time series cash flow forecasting framework provides a methodological foundation for moving from passive prediction to proactive decision making. Its impact is not limited to accuracy improvement. It lies in tightly linking forecasting with strategy evaluation, risk control, and resource allocation. Looking ahead, broader cross-system data integration is expected to become standard practice. Transaction records, invoices, and contracts, supply chain execution, financing terms, and interest rate information can be incorporated into a unified temporal representation space. This would further enhance the coverage and interpretability of cash flow drivers. Under continued advances in financial technology and digital corporate governance, such explainable, auditable, and boundary-aware forecasting frameworks are likely to become core components of fund management and risk governance. They can support more resilient operational decisions and higher-quality risk management in uncertain environments.

References

[1] J. Yan and H. Wang, "Self-interpretable time series prediction with counterfactual explanations," International Conference on Machine Learning, PMLR, 2023.

[2] E. De Brouwer, "Deep counterfactual estimation with categorical background variables," Advances in Neural Information Processing Systems, vol. 35, pp. 35213-35225, 2022.

[3] M. Z. Hossan and T. Sultana, "Causal Inference in Business Decision-Making: Integrating Machine Learning with Econometric Models for Accurate Business Forecasts," International Journal of Technology, Management and Humanities, vol. 9, no. 1, pp. 11-24, 2023.

[4] T. Crasson, Y. Nabet and M. Lécuyer, "Training and Evaluating Causal Forecasting Models for Time-Series," arXiv preprint arXiv:2411.00126, 2024.

[5] A. Yamaguchi et al., "Learning counterfactual explanations with intervals for time-series classification," Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, 2024.

[6] Y. Lu, "Policy-Aware Enterprise Risk Assessment through LLM-Based Cross-Source Representation and Reasoning," 2024.

[7] T. Karatas, F. Klinkert and A. Hirsa, "Supervised Neural Networks for Illiquid Alternative Asset Cash Flow Forecasting," arXiv preprint arXiv:2108.02853, 2021.

[8] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, vol. 9, no. 8, pp. 1735-1780, 1997.

[9] S. Siami-Namini, N. Tavakoli and A. S. Namin, "The performance of LSTM and BiLSTM in forecasting time series," 2019 IEEE International Conference on Big Data (Big Data), IEEE, 2019.

[10] A. Vaswani et al., "Attention is all you need," Advances in Neural Information Processing Systems, vol. 30, 2017.

[11] H. Zhou et al., "Informer: Beyond efficient transformer for long sequence time-series forecasting," Proceedings of the AAAI Conference on Artificial Intelligence, vol. 35, no. 12, 2021.

[12] Y. Liu et al., "iTransformer: Inverted transformers are effective for time series forecasting," arXiv preprint arXiv:2310.06625, 2023.

[13] S. Wang et al., "TimeMixer: Decomposable multiscale mixing for time series forecasting," arXiv preprint arXiv:2405.14616, 2024.